arXiv:cs/0407016v1 [cs.AI] 6 Jul 2004

Learning for Adaptive Real-time Search

Vadim Bulitko

Department of Computing Science, University of Alberta,
Edmonton, Alberta T6G 2E8, Canada
bulitko@ualberta.ca
http://www.cs.ualberta.ca/lrts

Abstract. Real-time heuristic search is a popular model of acting and learning in intelligent autonomous agents. Learning real-time search agents improve their performance over time by acquiring and refining a value function guiding the application of their actions. As computing the perfect value function is typically intractable, a heuristic approximation is acquired instead. Most studies of learning in real-time search (and reinforcement learning) assume that a simple value-function-greedy policy is used to select actions. This is in contrast to practice, where high-performance is usually attained by interleaving planning and acting via a lookahead search of a non-trivial depth. In this paper, we take a step toward bridging this gap and propose a novel algorithm that (i) learns a heuristic function to be used specifically with a lookahead-based policy, (ii) selects the lookahead depth adaptively in each state, (iii) gives the user control over the trade-off between exploration and exploitation. We extensively evaluate the algorithm in the sliding tile puzzle testbed comparing it to the classical LRTA* and the more recent weighted LRTA*, bounded LRTA*, and FALCONS. Improvements of 5 to 30 folds in convergence speed are observed.

Keywords: real-time heuristic search, planning and learning, on-line learning, adaptive lookahead search.

1 Problem Formulation

Complete search methods such as A* and IDA* [1] produce optimal solutions with admissible heuristics. The price of optimality is the substantial running time often exponential in the problem dimension [2]. This limits the applicability of complete search in tasks with large state spaces and limited time per action. The body of research on real-time search trades off optimality of solution for running time [3]. Consequently, these techniques are widely applied to real-time path-planning, game-playing, control, and general decision-making. Most approaches to building boundedly rational [4] real-time decision-making agents interleave lookahead-based deliberation and backing up the information to select an action [3,4,5]. Accordingly, most efforts to increase the rationality of such agents fall into three categories: (i) better hand-engineered and automatically derived heuristic functions [6], (ii) various lookahead control, state pruning methods, and search extension techniques [4], and (iii) specialized hardware [7].

In this paper we will focus on the first two ways of increasing rationality of autonomous decision-making agents. Namely, we will consider the framework of *learning* in real-time heuristic search [3,8] henceforth referred to as LRTS. It is an attractive model of decision-making in autonomous agents since it allows the agent to improve its performance over repeated trials in the same environment. Hence, another way to view LRTS is in the light of on-line reinforcement learning (RL) [5]. Not only the learning ability enables a gradual improvement of the solution, but also it allows the agent to

act with an incomplete model of the environment, within non-stationary environment and goals, and in a non-deterministic environment. Consequently, the LRTS model has been successfully applied to numerous practical tasks including moving target search problems [9], robot navigation and localization [10], and robot exploration [11]. Several important attributes of LRTS algorithms and RL agents are:

• **final solution quality** is measured in relation to the optimal solution quality. Having upper bounds on the cost of the solution the agent eventually converges to is important for performance guarantees;

• **convergence speed** is measured in the total number of actions the agent executes before it converges to its final solution;

• **resource bounds** are imposed on the amount of memory the agent requires to converge to its final solution. This can be crucial for autonomous agents as their hardware is particularly limited;

• **exploration versus exploitation control.** As finding high-quality solutions requires an extensive exploration of the state space, there is usually a trade-off between the final solution quality and the speed of the convergence and the resources required. Thus, an important attribute of an LRTS algorithm is whether this trade-off can be user-controlled;

• **convergence stability** stems from a related exploration and exploitation trade-off. Namely, fast convergence to better solutions requires an aggressive exploration of the state space. Such an LRTS agent usually demonstrates "optimism in the face of uncertainty" [12] by optimistically abandoning already found solution in an eager exploration of unknown regions of the state space. As a result, the solution quality can vary by orders of magnitude in consecutive trials. As argued in [13], this may be unacceptable depending on the application;

• **integration of learning and planning** is critical in most real-world LRTS (and learning game-playing) agents as their high performance crucially depends on extensive lookahead that constitutes the planning step of each cycle [7,14,15]. Contrary to a popular belief, deeper lookahead does not necessarily improve the decision-making quality [16,17,18]. Thus, an adaptive control of lookahead (i.e., planning) [4] and its integration with the heuristic function update (i.e., learning) appear promising.

2 Related Research

We will now consider several LRTS algorithms in light of the attributes discussed in the previous section. LRTA* [3] is an early and still widely used LRTS algorithm. Under certain assumptions, it is guaranteed to convergence to an optimal solution in a finite number of trials. Being an optimality-seeking algorithm, LRTA* can require prohibitive computational resources (memory), is unstable in convergence [13], does not provide an exploration vs. exploitation control, and does not learn the heuristic function tailored to its lookahead module (called 'mini-min' by Korf).

Several extensions of LRTA* have been proposed. Weighted LRTA* [13] is a combination of LRTA* with an inadmissible initial heuristic function. The heuristic is required to be within $(1 + \epsilon)$ factor of the optimal heuristic function (i.e., ϵ-admissible). The ϵ parameter controls the suboptimality of the final solutions as weighted LRTA* converges to a solution with the cost within $(1 + \epsilon)$ of optimal. Consequently, larger values of ϵ lead to faster convergence and smaller memory requirements. The update

rule is taken directly from LRTA* and does not take the lookahead into account. As in LRTA*, convergence stability is not directly addressed.

Bounded LRTA* [13] uses additional memory to maintain an upper bound on the true heuristic function in each explored state. There is a user-set parameter δ used together with the upper bound to control the amount of exploration thereby increasing convergence stability. On the other hand, bounded LRTA* incurs an additional memory and running time overhead as well as additional complexity of acquiring a non-trivial upper bound and refining it during the search. Additionally, its optimality seeking nature (for $\delta > 2$) can lead to intractable storage requirements. Consequently, a combination of weighted and bounded LRTA* extensions is needed. The resulting interplay between ϵ and δ parameters and its effects on the algorithm properties is a subject of future research [13, p. 35]. Finally, there is still no explicit consideration of lookahead in the learning module of bounded LRTA* and it is not clear how well the weighting and upper-bounding techniques will work with the lookahead search of a non-trivial depth.

FALCONS [19] maintains an additional heuristic function which is a lower bound on the cost between the current state and the initial state. It was designed to speed up the convergence while retaining the optimality of the final solution and therefore can require prohibitive storage. Additionally, it does not offer the user any control over the exploration vs. exploitation trade-off or convergence stability. Finally, FALCONS does not make specific considerations of the lookahead in its learning module.

3 Novel Learning Real-time Search Algorithm

Intuition. While past efforts have often focused *either* on gaining closer approximations to the true distance to goal [3,6] *or* better management of the planning phase [4], we recognize that these two are tightly coupled as the heuristic function is useful only insomuch as a guidance to the decision-making module of the agent. Consequently, instead of learning or manually engineering heuristic functions that approximate the true distance to goal well, we design the learning module specifically for the adaptive lookahead search. Below is the intuition behind the novel LRTS algorithm:

1. Since optimal solutions are inherently intractable in many non-trivial domains (e.g., [2]), we allow our algorithm to converge to suboptimal solutions. This results in significant speed-ups in convergence and savings in memory. At the same time, similar to weighted LRTA*, we upper-bound the converged solutions.
2. Deep lookahead is wasteful in the states where the optimal action can be determined with a shallow search. Thus, our algorithm adjusts the lookahead depth adaptively based on the concept of so-called traps introduced later in the paper. Additionally, we execute a (short) variable-length sequence of actions per one lookahead thereby reducing the running time.
3. We propose a tighter coupling between the heuristic function update and lookahead modules by updating the heuristic function only when the agent runs into a trap with respect to its lookahead module. This results in fewer heuristic function updates and smaller storage requirements.
4. Similar to weighted LRTA*, we provide user with control over the optimality of the solutions and the resources required. The parameter can be selected based on the application at hand.

Traps. A lookahead based real-time agent such as RTA* or LRTA* conducts lookahead of a fixed depth expanding a potentially large number of states at each step. Upon taking an action, it updates the heuristic function in its current state and repeats the process until the goal state is found. Effectively, it follows the gradient of its heuristic function until it reaches a local minimum. This can be visualized as a 'pit' in the heuristic function surface where the heuristic values of *all* surrounding states together with the distances to them exceed the heuristic value of the current state. An LRTA*-like agent will continue moving inside the 'pit' and raising the heuristic values until the 'pit' is completely 'filled'. Such a 'filling process' can take a large number of actions and updates to the heuristic function. We formalize such 'pits' with a more general concept of γ-traps and incorporate explicit and efficient trap detection and recovery methods into our LRTS algorithm.

Search problem. Before we present the algorithm, the search task needs to be formally defined. An LRTS problem is defined as a tuple (S, A, c, s_0, S_g) where S is a finite set of states; A is a finite set of actions; $c : S \times A \to (0, \infty)$ is the cost function with $c(s, a)$ being the incurred cost of taking action a in state s; s_0 is the initial state; and $S_g \subset S$ is the set of goal states. We adopt the assumptions of [13] that every action is reversible in every state, every applicable action leads to another state, and at least one goal state in S_g is reachable from s_0.

Then the true cost of traveling from state s_1 to state s_2 is denoted by $\mathrm{dist}(s_1, s_2)$ and is defined as the minimal cumulative action cost the agent is to incur by traveling from s_1 to s_2. For any state s its true cost is defined as the minimal travel cost to the nearest goal: $h^*(s) = \min_{s_g \in S_g} \mathrm{dist}(s, s_g)$. A heuristic approximation h to the true cost is called admissible iff $\forall s \in S\,[h(s) \leq h^*(s)]$. The value of h in state s will be referred to as the heuristic cost of state s. Depth d child of state s is any state s' reachable from s in the minimum of d actions (denoted by $\|s, s'\| = d$). Depth d neighborhood of state s is then defined as $S_{s,d} = \{s_d \in S \mid \|s, s_d\| = d\}$. A state s is called a γ-trap of depth d iff it lacks a child of depth up to d such that the cost of getting to that child adjusted by γ plus the heuristic cost of the child does not exceed the heuristic cost of the current state: $\neg\left[\exists d' \in \overline{1 \ldots d}\; \exists s' \in S_{s,d'}\left[\gamma \cdot \mathrm{dist}(s, s') + h(s') \leq h(s)\right]\right]$.

The algorithm. Figure 1 presents our LRTS algorithm called γ-Trap as a control policy outputting a series of actions $\boldsymbol{a}$ in the state s. Thus, for a given search problem (S, A, c, s_0, S_g) the agent's current s state is initialized to s_0 and the heuristic function h is initialized to an *admissible* initial function h_0. The γ-Trap policy is then called repeatedly by the environment and the series of actions $\boldsymbol{a}$ it returns are used to update the current state until a goal state is reached. During this process, the heuristic function h is updated by the policy as well. Once a goal state is reached, the current problem-solving trial is deemed completed and the current state is reset to its initial location s_0. The agent is then to solve the problem again starting with the heuristic function h from the previous problem-solving trial. The convergence is achieved when a problem-solving trial is completed without any updates to the heuristic function h. The solution produced (i.e., the sequence of actions the agent took from s_0 to the goal state) is considered to be the final solution. The cumulative cost of its actions is the final solution cost.

γ-Trap operates as follows. Lines 1 through 5 conduct the lookahead search by incrementally expanding the neighborhood of the current state s up to the maximum

γ-**Trap**$(s, h, d_{\max}, \gamma)$

INPUT:

s: the current state
h: heuristic function
$d_{\max}$: maximum lookahead depth
γ: solution quality control parameter

OUTPUT: a series of actions $\boldsymbol{a}$

1 **for** $d = 1$ to $d_{\max}$ do
2 generate depth d neighborhood $S_{s,d} = \{s_d \in S \mid \|s, s_d\| = d\}$
3 **if** s is *not* a γ-trap of depth d: i.e., $\exists s' \in S_{s,d}\, [\gamma \cdot \mathrm{dist}(s, s') + h(s') \leq h(s)]$
4 **return** actions $\boldsymbol{a}$ leading from s to $\arg \min_{s_d \in S_{s,d}} (\gamma \cdot \mathrm{dist}(s, s_d) + h(s_d))$
5 **end for loop**
6 update the heuristic function: $h(s) \leftarrow \max_{d=1,\ldots,d_{\max}} \min_{s_d \in S_{s,d}} (\gamma \cdot \mathrm{dist}(s, s_d) + h(s_d))$
7 **return** actions $\boldsymbol{a}$ leading from s to the previous current state

Fig. 1. LRTS algorithm γ-Trap expressed as agent control policy.

depth $d_{\max}$ specified by the user. At each iteration of the expansion process if the current state s is found to be a *non*-trap (line 3) then the lookahead search is terminated and the *series* of actions $\boldsymbol{a}$ leading to the most promising state within the last expanded neighborhood is returned (line 4). Ties are resolved randomly. Note that this operation exits the policy call. The agent will then move to the new current state s. The heuristic function h will *not* be updated. On the other hand, if the loop terminates naturally in line 6 then the current state s is a γ-trap of depth $d_{\max}$. In this case the heuristic function needs to be updated. In order to minimize (i) the number of updates and (ii) the number of actions the agent takes while 'filling the pit', we do two things. First, we increase the heuristic $h(s)$ by the maximum allowable amount in line 6 (hence $\max$ over d-neigborhoods). Second, we backtrack to the previous state (line 7) and avoid wasteful wandering inside the 'pit'. Note that if the current state s is the initial state s_0 no backtracking is performed and the current state is not changed.

Completeness & final solution quality. Despite giving up final solution optimality, γ-Trap is complete and converges to a reasonable final solution. More precisely, the amount of optimality loss can be upper-bounded:

Theorem. *For any search problem (S, A, c, s_0, S_g) that satisfies the aforementioned conditions and given an admissible initial heuristic h_0, γ-Trap will find a finite solution on every trial. Additionally, over a finite number of repeated trials, γ-Trap will converge to a final solution of the cost upper-bounded by $\frac{h^*(s_0)}{\gamma \min_{S \times A} c(s,a)}$.*

The proof is available on our website and cannot be reproduced here due to the space constraints. Also note that the actual performance of γ-Trap is substantially better than the upper bound suggests. For instance, in the convergence experiments detailed below the average final solution cost for $\gamma = 0.2$ was only 144% of the optimal which is substantially below the upper bound of 500% suggested by the theorem.

4 Empirical Evaluation

We evaluated the performance of γ-Trap against existing algorithms: weighted LRTA* due to Shimbo and Ishida [13] and LRTA* and RTA* due to Korf [3] which we had re-

implemented. Additionally, since we extended the experimental setup used in [13], the figures for FALCONS and bounded LRTA* reported there are used in the comparisons as well.[1] Our extensions to the experiments reported in [13] were as follows: (i) we used up to 10 folds in order to measure the standard deviation and (ii) the algorithms were run with the lookahead depth fixed at 2, 5, 10, and 15 plies in addition to the lookahead of one reported therein. Due to the extensive nature of the results and the space constraints of the paper, only a subset of the findings will be reported herein. Full results can be found in a technical report on our web site.

LRTA* and RTA* were reimplemented directly from [3]. Additionally, a version of LRTA* upgraded with full hash-based lookahead pruning, consistent use of table-based h function, forced monotonicity of $h + g$ values, and only upward updates to h (as in [13]) was implemented and run under the name of iLRTA*. It was used for weighted LRTA* experiments and is listed as eiLRTA-x.x in the figures below where x.x is the value of ϵ. Likewise, bounded LRTA* is listed as biLRTA-x.x in the graphs (x.x being the value of δ). In order to investigate the impact of backtracking in γ-Trap, we also implemented a version of the algorithm with no backtracking upon h-function update. Instead, the algorithm always moves into the child node with the minimum $h + \gamma \operatorname{dist}$ value. In the plots, the backtracking version of γ-Trap is listed as gTrap-BT-x.x whereas the non-backtracking version is listed as gTrap-x.x (x.x is the value of γ). Manhattan distance was used as the initial heuristic function h_0 for all algorithms except weighted LRTA* where we used Manhattan distance multiplied by $(1 + \epsilon)$.

Convergence Speed. Within each fold, a random set of one hundred 8-puzzles was generated. Each of the algorithms (with the exception of non-learning RTA*) was run repeatedly on each of the puzzle instances until a convergence was achieved (indicated by the lack of h-function updates between the start and the goal states). The convergence cost and the final solution cost relative to the optimal were averaged over 100 puzzles. The the mean and standard deviation were computed over 10 folds of 100 puzzles each. The experiments were repeated for the lookahead depths of 1, 2, 5, 10, and 15. The results for the lookahead of one are shown in Figure 2. As the data indicate, γ-Trap with backtracking (gTrap-BT) outperforms all other algorithms in terms of the convergence cost: approximately seven times better than the most successful weighted LRTA*, 13 times better than FALCONS, and 31 times better than classic LRTA*. At the same time, it is comparable to others in terms of the final solution cost: 144% of optimal vs. 120% in the case of the most successful weighted LRTA*. This trend was observed for all lookahead depths.

Memory Requirements. We paralleled the experiments reported in [13] and computed the number of 15-puzzles from the one hundred puzzle set published by Korf in [1] that can be solved until convergence within four million heuristic function values stored. The results can be found in Figure 3 for the lookahead of one. In the experiments, γ-Trap was the only algorithm converging on all 100 Korf's puzzles within four million nodes. At the same time, it has found the highest quality solutions among the algorithms (e.g., 110% of optimal vs. 520% of optimal with the most successful weighted LRTA*).

[1] NB: some of the figures reported in [13] are incorrect due to a bug in their code. We thus use the figures recently re-computed by Shimbo [20].

Fig. 2. Performance over repeated trials with the lookahead of one. Each number is averaged over 10 independent folds of 100 random 8-puzzle instances each. Each 8-puzzle was solved until convergence. The total costs are plotted in the top graph while the final solution costs relative to optimal are below. The data for bounded LRTA* and FALCONS are taken from [13] and computed over only one fold (hence no error bars or final solution costs).

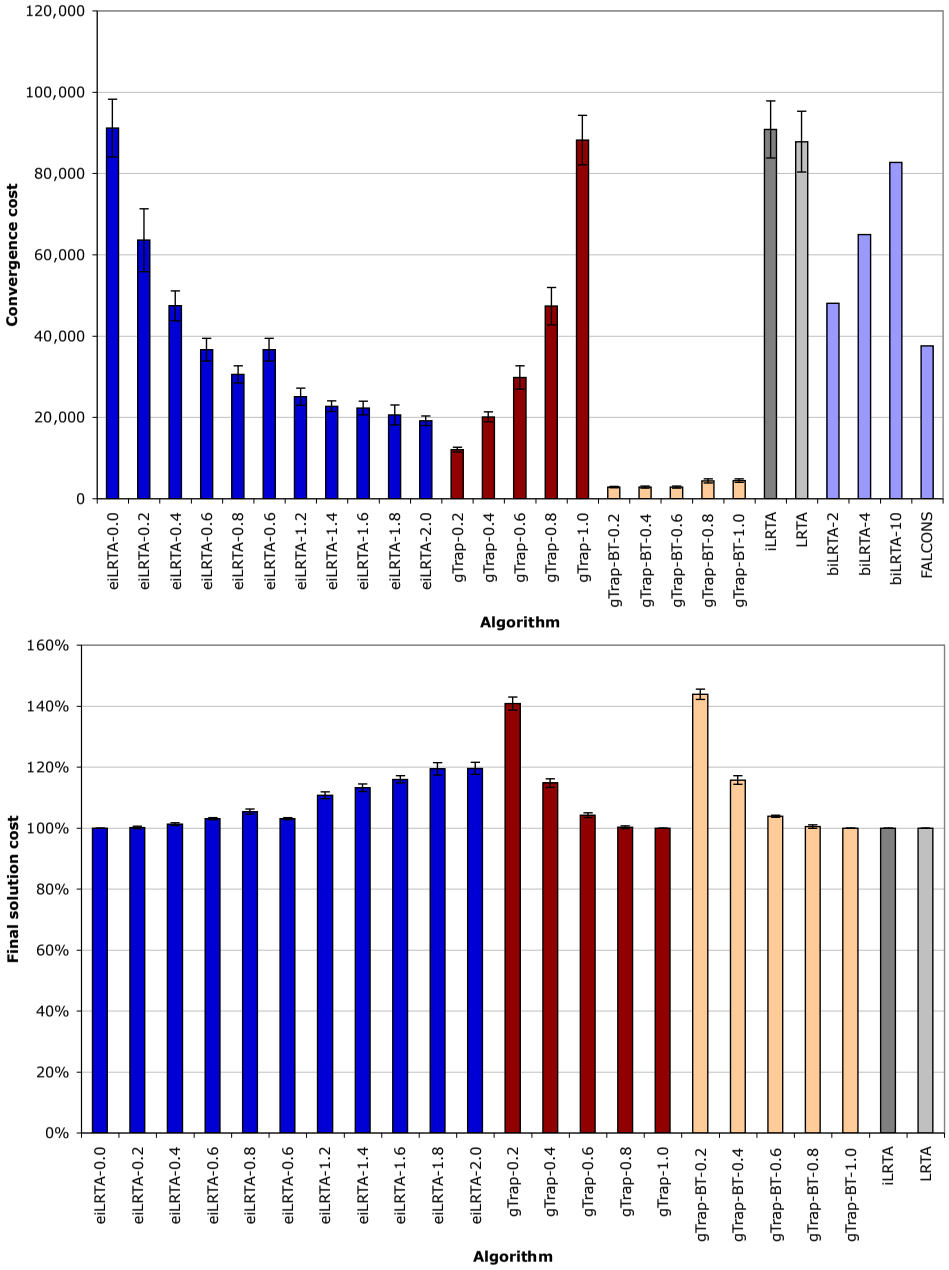

Fig. 3. Top: the number of Korf's 15-puzzles solved to convergence within four million stored heuristic values. Bottom: the final solution quality attained. The lookahead is fixed at one.

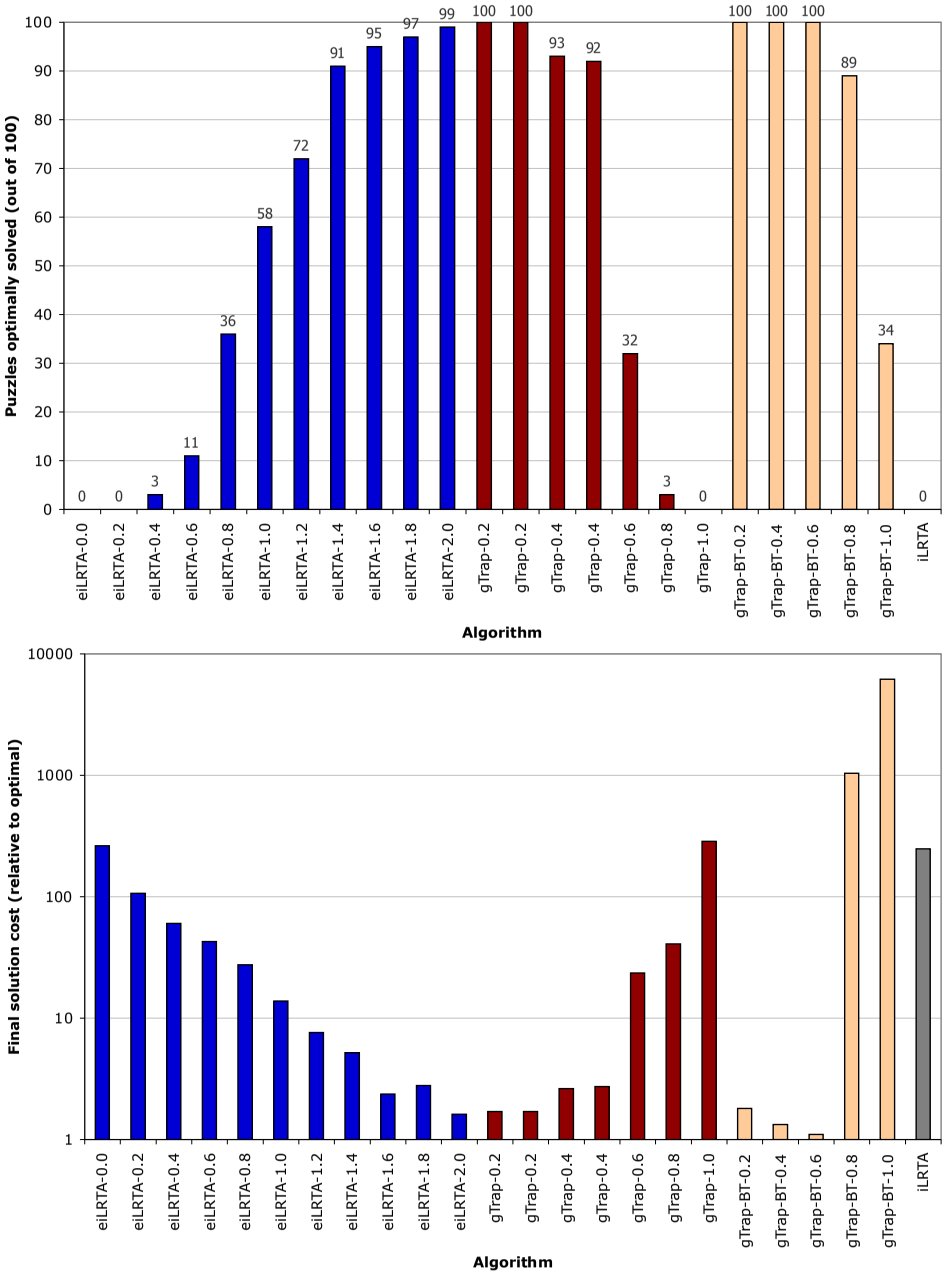

Convergence Stability. Even though γ-Trap was not designed to minimize the oscillations in the solution cost during the convergence process, we evaluated it against the standard algorithms along that dimension. Paralleling the study of [13], we adopted IAE, ISE, ITAE, ITSE, and SOD stability indices. IAE and SOD reported in this paper are defined as: IAE $= \text{avg}_s \sum_{i=1}^{\infty} |\,\text{cost}(s,i) - h^*(s)|$, SOD $= \text{avg}_s \sum_{i=1}^{\infty} \max\left(0, \text{cost}(s,i+1) - \text{cost}(s,i)\right)$ where $\text{cost}(s,i)$ is the state s's solution cost on ith trial while $h^*(s)$ is the optimal cost. The 8-puzzle results for IAE and SOD with the lookahead of one are plotted in Figure 4. Similar improvements were observed for deeper lookaheads. The data for bounded LRTA* and FALCONs are taken from [13] with the proper scaling factors of: 10^4 for IAE, 10^7 for ISE, 10^6 for ITAE, 10^9 for ITSE, and 10^3 for SOD [20]. We observe that γ-Trap with backtracking learns in a significantly more stable fashion than all other algorithms including bounded LRTA* specifically designed for stable convergence. In particular, for $\gamma = 0.2$ it breaks the past records by nearly 5 folds (in SOD) and over 14 folds (in IAE). γ-Trap without backtracking appears comparable to weighted LRTA*.

First-trial Performance. We evaluated the first-trial performance of the LRTS algorithms on random sets of 8-puzzles and Korf's set of 15-puzzles. Namely, within a single fold, each algorithm was presented with 100 puzzles each of which it had to solve within 500 thousand moves. The experiments were repeated over 10 independent folds each with the lookahead of 1, 2, 5, 10, and 15. The results are plotted in Figure 5. We included the data reported by Shimbo and Ishida in [13,20] in the bottom graph. The data suggest that the superior convergence speed, memory requirements, and learning stability of γ-Trap come at the cost of a more expensive first-trial solution. This behavior is similar to that exhibited by FALCONS. Additionally, we note that the backtracking part of γ-Trap appears to be the primary factor of its superior learning and inferior first-trial performance as the no-backtracking version (gTrap) is similar to the weighted LRTA* in both respects.

5 Future Work & Conclusions

Learning real-time search is an appealing model for the study of rational autonomous decision-making agents. The real-time nature allows to test various strategies of interleaving planning and acting while the learning feature enables the agent to cope with unknown, uncertain, and non-stationary environments. Complete search methods, such as IDA*, face the intractability problem apparent even in toy domains (e.g., the scalable sliding tile puzzle). Real-time search agents trade intractable optimal solutions for effectively computable satisfactory solutions. In this paper we discuss important attributes of a learning real-time search agent and the associated trade-offs. We then examine the existing methods including LRTA* and its extensions as well as FALCONS in the light of the desired properties. A novel algorithm, called γ-Trap, is then founded on a tighter integration of adaptive lookahead-based planning and learning modules. Under the standard assumptions, it is proved to be complete and enjoys convergence to satisfactory solutions. We then evaluate it empirically against the existing methods and show a significant improvement in the convergence speed and stability as well as much reduced memory requirements. Future research directions include: deriving tighter upper bounds on convergence speed, stability, and final solution cost, controlling the amount

Fig. 4. Learning stability over 10 folds of 100 hundred random 8-puzzles each with the lookahead of one. SOD (top) and IAE (bottom) stability indices are plotted. The data for bounded LRTA* and FALCONs are from [13], were collected over a single fold, and thus do not have standard deviations plotted as the error bars.

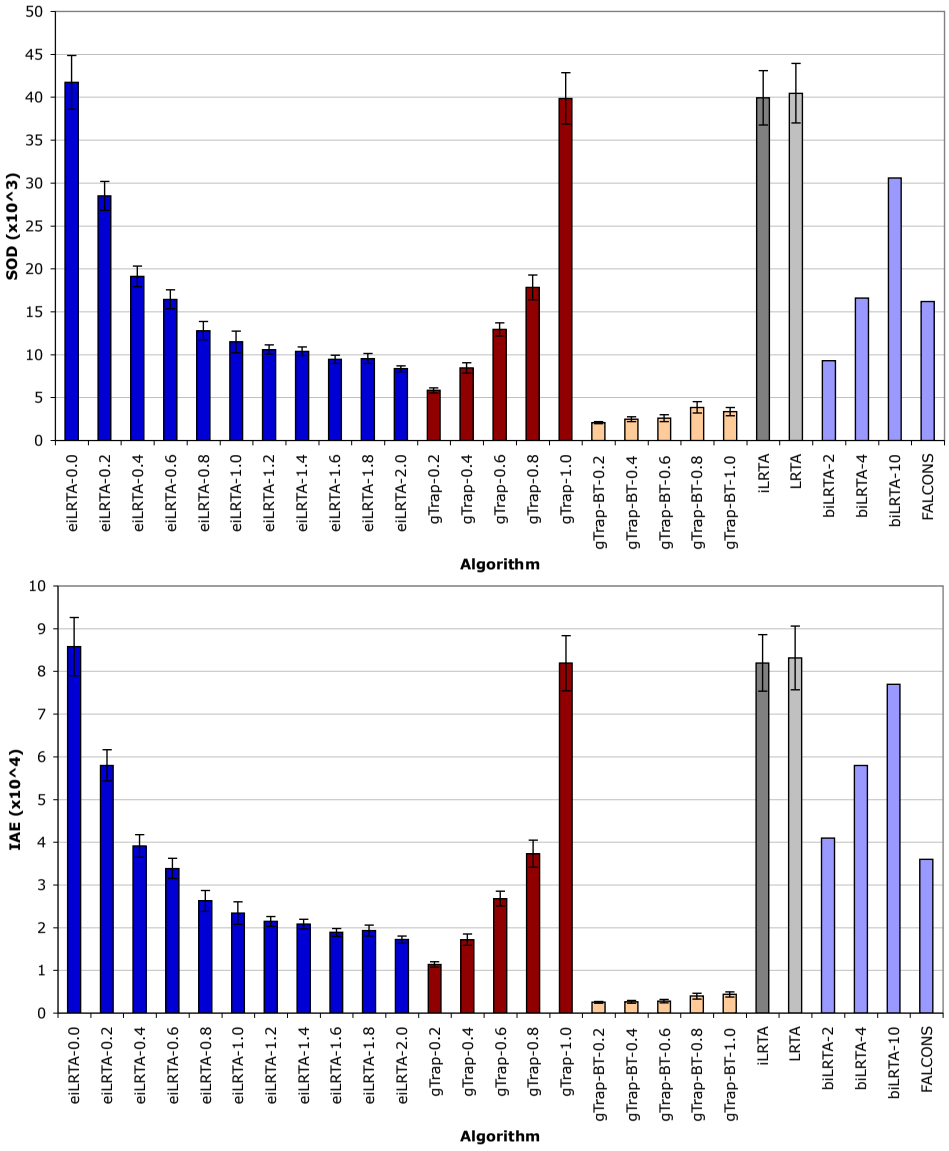

Fig. 5. First-trial solution quality with the lookahead of one. The top graph is for 10 folds of 100 random 8-puzzles each. The bottom graph is for one fold of Korf's one hundred 15-puzzles.

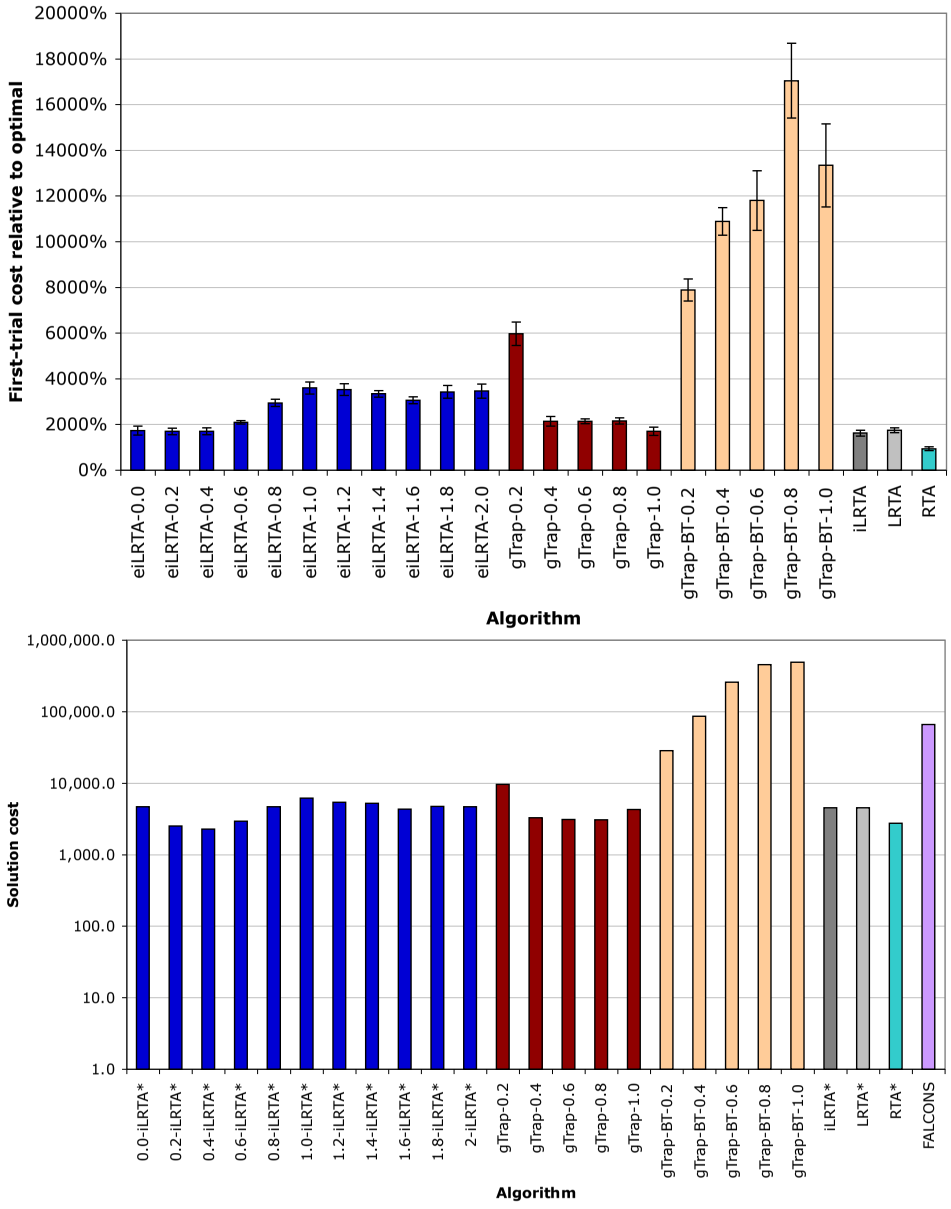

of exploration on the first trial, automatic selection of the γ parameter, and applying γ-Trap in non-deterministic, changing, and unknown environments.

Acknowledgements

Contributions of Valeriy Bulitko, David Furcy, Masashi Shimbo, Sven Koenig, Rob Holte, Rich Korf, and Stuart Russell are much appreciated. We are grateful for the support from NSERC, the University of Alberta, and the Alberta Ingenuity Centre for Machine Learning.

References

1. Korf, R.: Depth-first iterative deepening : An optimal admissible tree search. Artificial Intelligence **27** (1985)
2. Ratner, D., Warmuth, M.: Finding a shortest solution for the $N \times N$ extension of the 15-puzzle is intractable. In: Proceedings of AAAI-86. (1986) 168–172
3. Korf, R.: Real-time heuristic search. Artificial Intelligence **42** (1990) 189–211
4. Russell, S., Wefald, E.: Do the right thing: Studies in limited rationality. MIT Press (1991)
5. Barto, A.G., Bradtke, S.J., Singh, S.P.: Learning to act using real-time dynamic programming. Artificial Intelligence **72** (1995) 81–138
6. Korf, R.E., Taylor, L.A.: Finding optimal solutions to the twenty-four puzzle. In: Proceedings of the Thirteenth National Conference on Artificial Intelligence (AAAI-96), Portland, Oregon, USA, AAAI Press / The MIT Press (1996) 1202–1207
7. Hsu, F., Campbell, M., Hoane, A.: Deep Blue system overview. In: Proceedings of the 9th ACM Int. Conf. on Supercomputing. (1995) 240–244
8. Ishida, T.: Real-time Search for Learning Autonomous Agents. Kluwer Academic Publishers (1997)
9. Ishida, T., Korf, R.: Moving target search. In: Proceedings of the International Joint Conference on Artificial Intelligence. (1991) 204–210
10. Koenig, S., Simmons, R.: Solving robot navigation problems with initial pose uncertainty using real-time heuristic search. In: Proceedings of the International Conference on Artificial Intelligence Planning Systems. (1998)
11. Koenig, S.: Exploring unknown environments with real-time search or reinforcement learning. In: Proceedings of the Neural Information Processing Systems. (1999) 1003–1009
12. Moore, A., Atkeson, C.: Prioritized sweeping: Reinforcement learning with less data and less time. Machine Learning **13** (1993) 103–130
13. Shimbo, M., Ishida, T.: Controlling the learning process of real-time heuristic search. Artificial Intelligence **146** (2003) 1–41
14. Schaeffer, J., Culberson, J., Treloar, N., Knight, B., Lu, P., Szafron, D.: A world championship caliber checkers program. Artificial Intelligence **53** (1992) 273–290
15. Buro, M.: Probcut: An effective selective extension of the alpha-beta algorithm. ICCA Journal **18** (1995) 71–76
16. Nau, D.: Pathology on game trees revisited, and an alternative to minimaxing. Artificial Intelligence **21** (1983) 221–244
17. Bulitko, V., Li, L., Greiner, R., Levner, I.: Lookahead pathologies for single agent search. In: Proceedings of International Joint Conference on Artificial Intelligence. (2003) 1531–1533
18. Bulitko, V.: Lookahead pathologies and meta-level control in real-time heuristic search. In: Proceedings of the 15th Euromicro Conference on Real-Time Systems. (2003) 13–16
19. Furcy, D., Koenig, S.: Speeding up the convergence of real-time search. In: Proceedings of the National Conference on Artificial Intelligence. (2000) 891–897
20. Shimbo, M.: Personal communication. (2004)